arXiv:1602.04031v2 [math.CO] 11 May 2016

Proceedings of the 27th International Conference on Probabilistic, Combinatorial and Asymptotic Methods for the Analysis of Algorithms
Kraków, Poland, 4-8 July 2016

Counting Zeros in Random Walks on the Integers and Analysis of Optimal Dual-Pivot Quicksort

Martin Aumüller[1], Martin Dietzfelbinger[2], Clemens Heuberger[3‡], Daniel Krenn[4‡], Helmut Prodinger[5§]

[1] *IT University of Copenhagen, Denmark, maau@itu.dk*
[2] *Institut für Theoretische Informatik, Technische Universität Ilmenau, Germany, martin.dietzfelbinger@tu-ilmenau.de*
[3] *Institut für Mathematik, Alpen-Adria-Universität Klagenfurt, Austria, clemens.heuberger@aau.at*
[4] *Institut für Mathematik, Alpen-Adria-Universität Klagenfurt, Austria, math@danielkrenn.at or daniel.krenn@aau.at*
[5] *Department of Mathematical Sciences, Stellenbosch University, South Africa, hproding@sun.ac.za*

We present an average case analysis of two variants of dual-pivot quicksort, one with a non-algorithmic comparison-optimal partitioning strategy, the other with a closely related algorithmic strategy. For both we calculate the expected number of comparisons exactly as well as asymptotically, in particular, we provide exact expressions for the linear, logarithmic, and constant terms. An essential step is the analysis of zeros of lattice paths in a certain probability model. Along the way a combinatorial identity is proven.

Keywords: Dual-pivot quicksort, lattice paths, asymptotic enumeration, combinatorial identity

1 Introduction

Dual-pivot quicksort [Sed75, WNN15, AD15] is a family of sorting algorithms related to the well-known quicksort algorithm. In order to sort an input sequence $(a_1, \dots, a_n)$ of distinct elements, dual-pivot quicksort algorithms work as follows. (For simplicity we forbid repeated elements in the input.) If $n \le 1$, there is nothing to do. If $n \ge 2$, two input elements are selected as pivots. Let p be the smaller and q be the larger pivot. The next step is to partition the remaining elements into

- the elements smaller than p ("small elements"),
- the elements between p and q ("medium elements"), and
- the elements larger than q ("large elements").

‡C. Heuberger and D. Krenn are supported by the Austrian Science Fund (FWF): P 24644-N26 and by the Karl Popper Kolleg "Modeling–Simulation–Optimization" funded by the Alpen-Adria-Universität Klagenfurt and by the Carinthian Economic Promotion Fund (KWF).
§H. Prodinger is supported by an incentive grant of the National Research Foundation of South Africa.

Then the procedure is applied recursively to these three groups to complete the sorting.

The cost measure used in this work is the number of comparisons between elements. As is common, we will assume the input sequence is in random order, which means that each permutation of the n elements occurs with probability $1/n!$. With this assumption we may, without loss of generality, choose a_1 and a_n as the pivots. Even in this setting there are different dual-pivot quicksort algorithms; their difference lies in the way the partitioning is organized, which influences the partitioning cost. This is in contrast to standard quicksort with one pivot, where the partitioning strategy does not influence the cost—in partitioning always one comparison is needed per non-pivot element. In dual-pivot quicksort, the average cost (over all permutations) of partitioning and of sorting can be analyzed only when the partitioning strategy is fixed.

Only in 2009, Yaroslavskiy, Bentley, and Bloch [Yar09] described a dual-pivot quicksort algorithm that makes $1.9n \log n + O(n)$ comparisons [WNN15].[(i)] This beats the classical quicksort algorithm [Hoa62], which needs $2n \log n + O(n)$ comparisons on average. In [AD15], the first two authors of this article described the full design space for dual-pivot quicksort algorithms with respect to counting element comparisons. Among others, they studied two special partitioning strategies. The first one—we call it "Clairvoyant" in this work—assumes that the number of small and large elements is given (by an "oracle") before partitioning starts. It cannot be implemented; however, it is optimal among all partioning strategies that have access to such an oracle, and hence its cost provides a lower bound for the cost of all algorithmic partitioning strategies. In [AD15] it was shown that dual-pivot quicksort carries out $1.8n \log n + O(n)$ comparisons on average when this partitioning strategy is used. Further a closely related algorithmic partitioning strategy—called "Count" here—was described, which makes only $O(\log n)$ more comparisons on average than "Clairvoyant" and hence leads to a dual-quicksort variant with only $O(n)$ more comparisons.[(ii)]

One purpose of this paper is to make the expected number of comparisons in both variants precise and to determine the exact difference of the cost of these two strategies, both for partitioning and for the resulting dual-pivot quicksort variants.

Already in [AD15] it was noted that the exact value of the expected partitioning cost (i.e., the number of comparisons) of both strategies depends on the expected number of the zeros of certain lattice paths (Parts I and II). A complete understanding of this situation is the basis for our analysis of dual-pivot quicksort, which appears in Part III.

Lattice path enumeration has a long tradition. An early reference is [Moh79]; a recent survey paper is [Kra15]. As space is limited, many proofs and some additional results can be found in an appendix at arXiv:1602.04031v1.

2 Overview and Results

This work is split into three parts. We give a brief overview on the main results of each of these parts here. We use the Iversonian expression

$$[expr] = \begin{cases} 1 & \text{if } expr \text{ is true,} \\ 0 & \text{if } expr \text{ is false,} \end{cases}$$

(i) In this paper "log" denotes the natural logarithm to base e.

(ii) After completing this extended abstract we found a proof that "Count" is optimal among all algorithmic strategies. Details to be given in the full version.

popularized by Graham, Knuth, and Patashnik [GKP94].

The harmonic numbers and their variants will be denoted by

$$H_n = \sum_{m=1}^{n} \frac{1}{m}, \qquad H_n^{\text{odd}} = \sum_{m=1}^{n} \frac{[m \text{ odd}]}{m} \qquad \text{and} \qquad H_n^{\text{alt}} = \sum_{m=1}^{n} \frac{(-1)^m}{m}.$$

Of course, there are relations between these three definitions such as $H_n^{\text{alt}} = H_n - 2H_n^{\text{odd}}$ and $H_n^{\text{odd}} + H_{\lfloor n/2 \rfloor}/2 = H_n$, but it will turn out to be much more convenient to use all three notations.

Part I: Lattice Paths

In the first part we analyze certain lattice paths of a fixed length n. We start on the vertical axis, allow steps/increments $(1, +1)$ and $(1, -1)$ and end on the horizontal axis at $(n, 0)$. To be precise, the starting point on the vertical axis is chosen uniformly at random from the set $\{(0, -n), (0, -n + 2), \ldots, (0, n - 2), (0, n)\}$ of feasible points. Once this starting point is fixed, all paths to $(n, 0)$ are equally likely. We are interested in the number of zeros, denoted by the random variable Z_n, of such paths.

An exact formula for the expected number $\mathbb{E}(Z_n)$ of zeros is derived in two different ways (see identity (2.1) for these formulæ): On the one hand, we use the symbolic method and generating functions (see Appendix A), which gives the result in form of a double sum. This machinery extends well to higher moments and also allows us to obtain the distribution. The exact distribution is given in Appendix E; its limiting behavior as $n \to \infty$ is the discrete distribution

$$\mathbb{P}(Z_n = r) \sim \frac{1}{r(r+1)}.$$

On the other hand, a more probabilistic approach gives the expectation $\mathbb{E}(Z_n)$ as the simple single sum

$$\mathbb{E}(Z_n) = \sum_{m=1}^{n+1} \frac{[m \text{ odd}]}{m} = H_{n+1}^{\text{odd}},$$

see Section 4 for more details. The asymptotic behavior $\mathbb{E}(Z_n) \sim \frac{1}{2} \log n$ can be extracted (Appendix D).

The two approaches above give rise to the identity

$$\sum_{m=1}^{n+1} \frac{[m \text{ odd}]}{m} = \frac{4}{n+1} \sum_{0 \le k < \ell < \lceil n/2 \rceil} \frac{\binom{n}{k}}{\binom{n}{\ell}} + [n \text{ even}] \frac{1}{n+1} \left(\frac{2^n}{\binom{n}{n/2}} - 1 \right) + 1; \tag{2.1}$$

the double sum above equals the single sum of Theorem 4.1 by combinatorial considerations. One might ask about a direct proof of this identity. This can be achieved by methods related to hypergeometric sums and the computational proof is presented in Appendix C. We also provide a completely elementary proof which is "purely human".

Part II: More Lattice Paths and Zeros

The second part acts as connecting link between the lattice paths of fixed length of Part I and the dual-pivot quicksort algorithms of Part III.

The probabilistic model introduced in Section 3 (in Part I) is extended, and lattice paths are allowed to vary in length. For a number n (the number of elements to sort) the length of a path is the number of elements remaining when the two pivots, given by a random set of elements of size two, and the elements between these pivots are cut out.

The number of zeros X_n in this full model is analyzed; we provide again exact as well as asymptotic formulæ for the expectation $\mathbb{E}(X_n)$. Details are given in Section 7. Moreover, more specialized zero-configurations (needed for the analysis of different partitioning strategies in Part III) are considered as well (Section 6).

Part III: Dual-Pivot Quicksort

The main result of this work analyzes comparisons in the dual-pivot quicksort algorithm that uses the optimal (but unrealistic) partitioning strategy "Clairvoyant". Aumüller and Dietzfelbinger showed in [AD15] that this algorithm requires $1.8n \log n + O(n)$ comparisons on average, which improves on the average number of comparisons in quicksort ($2n \log n + O(n)$) and the recent dual-pivot algorithm of Yaroslavskiy et al. ($1.9n \log n + O(n)$, see [WNN15]). However, for real-world input sizes n the (usually negative) factor in the linear term has a great influence on the comparison count. Our asymptotic result is stated as the following theorem.

Theorem. *The average number of comparisons in the dual-pivot quicksort algorithm with a comparison-optimal partitioning strategy is*

$$\frac{9}{5} n \log n + An + B \log n + C + O(1/n)$$

as $n \to \infty$*, with* $A = \frac{9}{5}\gamma - \frac{1}{5}\log 2 - \frac{89}{25} = -2.659\ldots$.

The constants B and C are explicitly given, too, and more terms of the asymptotics are presented. The precise result is formulated as Corollary 10.2.

In fact, we even get an exact expression for the average comparison count. The precise result is formulated as Theorem 10.1. Moreover the same analysis is carried out for the partitioning strategy "Count", which is an algorithmic variant of the comparison-optimal strategy "Clairvoyant". Aumüller and Dietzfelbinger [AD15] could show that it requires $\frac{9}{5} n \log n + O(n)$ comparisons as well. In this paper we obtain the exact average comparison count (Theorem 10.3). The asymptotic result is again $\frac{9}{5} n \log n + An + O(\log n)$, but now with $A = -2.382\ldots$, so there is only a small gap between the average number of comparisons in the comparison-optimal strategy "Clairvoyant" and its algorithmic variant.

Part I: Lattice Paths

In this first part we analyze lattice paths of a fixed length n. These are introduced in Section 3 by a precise description of our probabilistic model. We will work with this model throughout Part I, and we analyze the number of zeros Z_n.

The outline is as follows: We derive an exact expression for the expected number $\mathbb{E}(Z_n)$ of zeros by the generating functions machinery in Appendix A; a more probabilistic approach can

be found in Section 4. Appendix D deals with asymptotic considerations. Direct proofs of the obtained identity are given in Appendix C and the distribution of Z_n is tackled in Appendix E.

3 Probabilistic Model

We consider paths of a given length n on the lattice $\mathbb{Z}^2$, where only steps $(1,+1)$ and $(1,-1)$ are allowed. These paths are chosen at random according to the rules below.

Let us fix a length $n \in \mathbb{N}_0$. A path P_n ending in $(n,0)$ (no choice for this end-point) is chosen according to the following rules.

1. First, choose a starting point $(0,S)$ where S is a random integer uniformly distributed in $\{-n,-n+2,\ldots,n-2,n\}$, i. e., $S=s$ occurs only for integers s with $|s|\le n$ and $s\equiv n \pmod 2$.

2. Second, a path is chosen uniformly at random among all paths from $(0,S)$ to $(n,0)$.

The conditions on S characterize those starting points from which $(n,0)$ is reachable.

We are interested in the number of intersections with the horizontal axis of a path. To make this precise, we define a *zero* of a path P_n as a point $(x,0)\in P_n$.

Thus, let P_n be a path of length n which is chosen according to the probabilistic model above and define the random variable

$$Z_n = \text{number of zeros of } P_n.$$

In the following sections, we determine the value of $\mathbb{E}(Z_n)$ exactly (Appendix A and Section 4), as well as asymptotically (Appendix D). In Appendix A, we use the machinery of generating functions. This machinery turns out to be overkill if we are just interested in the expectation $\mathbb{E}(Z_n)$. However, it easily allows extension to higher moments and the limiting distribution.

In Section 4, we follow a probabilistic approach, which first gives a result on the probability model that at the first glance looks surprising: the equidistribution at the initial values turns out to carry over to every fixed length of the remaining path. This result yields a simple expression for the expectation $\mathbb{E}(Z_n)$ in terms of harmonic numbers, and thus immediately yields a precise asymptotic expansion for $\mathbb{E}(Z_n)$. The generating function approach, however, gives the expectation in terms of a double sum of quotients of binomial coefficients (the right-hand side of (2.1)), see Appendix A.

Appendix C gives a direct computational proof that these two results coincide. The original expression in [AD15] (a double sum over a quotient of a product of binomial coefficients and a binomial coefficient) is also shown to be equal in Appendix C. Explicit as well as asymptotic expressions for the distribution $\mathbb{P}(Z_n=r)$ can be found in Appendix E.

4 A Probabilistic Approach

Theorem 4.1. *For a randomly (as described in Section 3) chosen path of length n, the expected number of zeros is*

$$\mathbb{E}(Z_n) = H^{\mathrm{odd}}_{n+1}.$$

Before proving the theorem, we consider an equivalent probability model for our random paths formulated as an urn model. A number R from $\{0, \dots, n\}$ is chosen uniformly at random. We place R red balls and $B = n - R$ black balls in an urn. Subsequently, in n rounds the balls are taken from the urn (without replacements), in each round choosing one uniformly at random. The color of the ball chosen in round i is denoted by U_i.

We construct a random walk $(W_i)_{0 \le i \le n}$ on $\{-n, \dots, n\}$ from $U_1, \dots, U_n$ by setting $W_0 = R - B = 2R - n$ and

$$W_i = \begin{cases} W_{i-1} + 1 & \text{if } U_i = \text{black}, \\ W_{i-1} - 1 & \text{if } U_i = \text{red} \end{cases}$$

for $1 \le i \le n$. In each step, W_i equals the difference of the number of remaining red and black balls in the urn. Clearly, then, $W_n = 0$.

One can look at the trajectories of this random walk, represented in the grid $\{0, \dots, n\} \times \{-n, \dots, n\}$ as sequences $((0, W_0), (1, W_1), \dots, (n, W_n))$. Appendix B explains the equivalence between the two models.

In order to prove Theorem 4.1, we need the following property of our paths.

Lemma 4.2. *Let $m \in \mathbb{N}_0$ with $m \le n$. The probability that a random path P_n (as defined in Section 3) runs through $(n - m, k)$ is*

$$\mathbb{P}((n - m, k) \in P_n) = \frac{1}{m + 1} \tag{4.1}$$

for all k with $|k| \le m$ and $k \equiv m \pmod 2$, otherwise 0.

The proof of this lemma can be found in Appendix B.

A closer look reveals that when we reverse the paths, our model is equivalent to a contagion Pólya urn model with two colors, starting with one ball of each color, where we sample with replacement and put another ball of the color just drawn into the urn. In this setting, uniform distribution for feasible points with the same first coordinate and hence the result of the lemma are well-known phenomena. These results and more on the urn model can be found, for example, in Mahmoud [Mah08].

We continue with the actual proof of our theorem.

Proof of Theorem 4.1: By Lemma 4.2, the expected number of zeros of P_n is

$$\mathbb{E}(Z_n) = \sum_{m=0}^{n} \mathbb{P}((n - m, 0) \in P_n) = \sum_{m=0}^{n} \frac{[m \text{ even}]}{m + 1} = \sum_{m=1}^{n+1} \frac{[m \text{ odd}]}{m} = H_{n+1}^{\text{odd}}. \qquad \square$$

5 Additional Results

The expected number of zeros can be evaluated asymptotically. We obtain

Corollary 5.1.

$$\mathbb{E}(Z_n) = \frac{1}{2} \log n + \frac{\gamma + \log 2}{2} + \frac{1 + [n \text{ even}]}{2n} - \frac{2 + 9[n \text{ even}]}{12n^2} + O\Big(\frac{1}{n^3}\Big)$$

asymptotically as n tends to infinity.

The proof of this result uses the well-known asymptotic expansion of the harmonic numbers. The actual asymptotic computations[iii] have been carried out using the asymptotic expansions module [HK15] of SageMath [Dev16], see Appendix D.

By combining the generating function and probabilistic approach we obtain the following identity.

Theorem 5.2. *For $n \geq 0$, we have*

$$\begin{aligned} &\frac{4}{n+1} \sum_{0 \leq k < \ell < \lceil n/2 \rceil} \frac{\binom{n}{k}}{\binom{n}{\ell}} + [n \text{ even}] \frac{1}{n+1} \left(\frac{2^n}{\binom{n}{n/2}} - 1 \right) + 1 \\ &= \frac{1}{n+1} \sum_{m=0}^{\lfloor n/2 \rfloor} \sum_{\ell=m}^{n-m} \frac{\binom{2m}{m}\binom{n-2m}{\ell-m}}{\binom{n}{\ell}} = H_{n+1}^{\mathrm{odd}}. \end{aligned}$$

The second expression for the expected number of zeros, but without taking the zero at $(n, 0)$ into account, has been given in [AD15, displayed equation after (14)]. In Appendix C we give two direct proofs of the identity above: One of them follows a computer generated proof ("creative telescoping") by extracting the essential recurrence. The second proof is "human" and completely elementary using not more than Vandermonde's convolution.

Furthermore, the generating function machinery allows us to determine the distribution of the number Z_n of zeros. Beside an exact formula (see Appendix E), we get the following asymptotic result.

Theorem 5.3. *Let $0 < \varepsilon \leq \frac{1}{2}$. For positive integers r with $r = O\big(n^{1/2-\varepsilon}\big)$, we have asymptotically*

$$\mathbb{P}(Z_n = r) = \frac{1}{r(r+1)} \big(1 + O\big(1/n^{2\varepsilon}\big)\big)$$

as n tends to infinity.

Part II: More Lattice Paths and Zeros

This second part deals with an analysis of some special zero-configurations, which are needed for the analysis of the partitioning strategies in Part III. Moreover, in Section 7, we extend the model introduced in Section 3 to accommodate lattice paths of variable length. Again expectations are studied exactly and asymptotically.

6 Going to Zero and Coming From Zero

For the analysis of comparison-optimal dual-pivot quicksort algorithms (see Part III) we need the following two variants of zeros on the lattice path.

- An *up-to-zero situation* is a point $(x, 0) \in P_n$ such that $(x-1, -1) \in P_n$.
- A *down-from-zero situation* is a point $(x, 0) \in P_n$ such that $(x+1, -1) \in P_n$.

[iii] A worksheet containing the computations can be found at `http://www.danielkrenn.at/downloads/quicksort-paths/quicksort-paths.ipynb`.

We show

$$\mathbb{E}(\text{number of up-to-zero situations on } P_n) = \frac{1}{2}\Big(\mathbb{E}(Z_n) - \frac{[n \text{ even}]}{n+1}\Big) = \frac{1}{2}H_n^{\text{odd}}$$

and

$$\mathbb{E}(\text{number of down-from-zero situations on } P_n) = \frac{1}{2}\left(\mathbb{E}(Z_n) - 1\right) = \frac{1}{2}\big(H_{n+1}^{\text{odd}} - 1\big).$$

Proof idea: The factor $\frac{1}{2}$ stems from symmetry: Up-to-zero situations at $(x, 0)$ occur with the same probability as the symmetric "down-to-zero" situations at $(x, 0)$, similarly for down-from-zero situations. The correction terms $\frac{[n \text{ even}]}{n+1}$ and 1 are caused by the fact that there is a zero, but no up-to-zero situation, at $(0, 0)$, and a zero, but no down-from-zero situation, at $(n, 0)$. The full proofs are in Appendix F. □

7 Lattice Paths of Variable Length

In this section, we use a random variable N' instead of the fixed n above. Let us fix an $n \in \mathbb{N}$ with $n \geq 2$. We choose a path length N' according to the following rules.

1. Choose (P, Q) with $1 \leq P < Q \leq n$ uniformly at random from all $\binom{n}{2}$ possibilities.
2. Let $N' = n - 1 - (Q - P)$.
3. Choose a path of length N' according to Section 3.

Let us denote the number of up-to-zero and down-from-zero situations on the path by $X_n^{\nearrow}$ and $X_n^{\searrow}$, respectively. In Appendix G, we show

$$\mathbb{E}\big(X_n^{\nearrow}\big) = \frac{1}{2\binom{n}{2}} \sum_{n'=0}^{n-2} \sum_{m=1}^{n'} [m \text{ odd}]\, \frac{n'+1}{m} = \frac{1}{2}H_{n-2}^{\text{odd}} - \frac{1}{8} + \frac{(-1)^n}{8(n - [n \text{ even}])}$$

and

$$\mathbb{E}\big(X_n^{\searrow}\big) = \frac{1}{2\binom{n}{2}} \sum_{n'=0}^{n-2} \sum_{m=3}^{n'+1} [m \text{ odd}]\, \frac{n'+1}{m} = \mathbb{E}\big(X_n^{\nearrow}\big) - \frac{1}{2} + \frac{1}{2(n - [n \text{ even}])}.$$

Part III: Dual-Pivot Quicksort

In this third and last part of this work, we finally analyze two different partitioning strategies and the dual-pivot quicksort algorithm itself.

As mentioned in the introduction, the number of comparisons of dual-pivot quicksort depends on the concrete partitioning procedure. For example, if one wants to classify a large element, i. e., an element larger than the larger pivot, comparing it with the larger pivot is unavoidable, but it depends on the partitioning procedure whether a comparison with the smaller pivot occurs, too.

First, in Section 8, we make our set-up precise, fix notions, and start solving the dual-pivot quicksort recurrence (8.1). This recurrence relates the cost of the partitioning step to the total number of comparisons of dual-pivot quicksort.

In Section 9 two partitioning strategies, called "Clairvoyant" and "Count", are introduced and their respective cost is analyzed. It will turn out that the results on lattice paths obtained in Parts I and II are central in determining the partitioning cost exactly.

Everything is put together in Section 10: We obtain the exact comparison count for two versions of dual-pivot quicksort (Theorems 10.1 and 10.3). The asymptotic behavior is extracted out of the exact results (Corollaries 10.2 and 10.4).

8 Solving the Dual-Pivot Quicksort Recurrence

We consider versions of dual-pivot quicksort that act as follows on an input sequence $(a_1, \dots, a_n)$ consisting of distinct numbers: If $n \le 1$, do nothing, otherwise choose a_1 and a_n as pivots, and by one comparison determine $p = \min(a_1, a_n)$ and $q = \max(a_1, a_n)$. Use a partitioning procedure to partition the remaining $n-2$ elements into the three classes *small*, *medium*, and *large*. Then call dual-pivot quicksort recursively on each of these three classes to finish the sorting, using the same partitioning procedure in all recursive calls.

Let P_n, a random variable, denote the *partitioning cost*. This is defined as the number of comparisons made by the partitioning procedure if the input $(a_1, \dots, a_n)$ is assumed to be in random order. Further, let C_n be the random variable that denotes the number of comparisons carried out when sorting n elements with dual-pivot quicksort. The reader should be aware that both P_n and C_n are determined by the partitioning procedure used.

Since the input $(a_1, \dots, a_n)$ is in random order and the partitioning procedure does nothing but compare elements with the two pivots, the inputs for the recursive calls are in random order as well, which implies that the distributions of P_n and C_n only depend on n. In particular we may assume that when the sorting algorithm is called on n elements during recursion, the input is a permutation of $\{1, \dots, n\}$.

The recurrence

$$\mathbb{E}(C_n) = \mathbb{E}(P_n) + \frac{3}{\binom{n}{2}} \sum_{k=1}^{n-2} (n-1-k)\, \mathbb{E}(C_k) \tag{8.1}$$

for $n \ge 0$ describes the connection between the expected sorting cost $\mathbb{E}(C_n)$ and the expected partitioning cost $\mathbb{E}(P_n)$. It will be central for our analysis. Note that it is irrelevant for (8.1) how the partitioning cost $\mathbb{E}(P_n)$ is determined; it need not even be referring to comparisons. The recurrence is simple and well-known; a version of it occurs already in Sedgewick's thesis [Sed75]. For the convenience of the reader we give a brief justification in Appendix H. In Hennequin [Hen91] recurrence (8.1) was solved exactly for $\mathbb{E}(P_n) = an + b$, where a and b are constants. For $\mathbb{E}(P_n) = an + O(n^{1-\varepsilon})$ the solution is $\mathbb{E}(C_n) = \frac{6}{5} an \log n + O(n)$, see [AD15, Theorem 1].

9 Partitioning Algorithms and Their Cost

In Section 8 we saw that in order to calculate the average number of comparisons $\mathbb{E}(C_n)$ of a dual-pivot quicksort algorithm we need the expected partitioning cost $\mathbb{E}(P_n)$ of the partitioning

procedure used. The aim of this section is to determine $\mathbb{E}(P_n)$ for two such partitioning procedures, "Clairvoyant" and "Count", to be described below.

We use the set-up described at the beginning of Section 8. For partitioning we use comparisons to *classify* the $n-2$ elements $a_2, \dots, a_{n-1}$ as *small*, *medium*, or *large*. We will be using the term *classification* for this central aspect of partitioning. Details of a partitioning procedure that concern how the classes are represented or elements are moved around may and will be ignored. (Nonetheless, in Appendix M we provide pseudocode for the considered classification strategies turned into dual-pivot quicksort algorithms.) The cost P_n depends on the concrete classification strategy used, the only relevant difference between classification strategies being whether the next element to be classified is compared with the smaller pivot p or the larger pivot q first. This decision may depend on the whole history of outcomes of previous comparisons. (The resulting abstract classification strategies may conveniently be described as classification trees, see [AD15], but we do not need this model here.)

Two comparisons are necessary for each medium element. Furthermore, one comparison with p is necessary for small and one comparison with q is necessary for large elements. As the input consists of the elements $1, \dots, n$, there are $p-1$ small, $q-p-1$ medium, and $n-q$ large elements. Averaging over all $\binom{n}{2}$ positions of the pivots, we see that on average

$$\frac{4}{3}(n-2)+1 \tag{9.1}$$

necessary comparisons are required no matter how the classification procedure works, see [AD15, (5)]; the summand $+1$ corresponds to the comparison of a_1 and a_n when choosing the two pivots.

We call other comparisons occurring during classification *additional comparisons*. That means, an additional comparison arises when a small element is compared with q first or a large element is compared with p first. In order to obtain $\mathbb{E}(P_n)$ for some classification strategy, we have to calculate the expected number of additional comparisons.

Next we describe two (closely related) classification strategies from [AD15]. Let s_i and ℓ_i denote the number of elements that have been classified as small and large, respectively, in the first i classification rounds. Set $s_0 = \ell_0 = 0$.

Strategy "Clairvoyant". *Assume the input contains $s = p-1$ small and $\ell = n-q$ large elements. When classifying the ith element, for $1 \le i \le n-2$, proceed as follows: If $s - s_{i-1} \ge \ell - \ell_{i-1}$, compare with p first, otherwise compare with q first.*

The number of additional comparisons of Clairvoyant is denoted by A_n^{cv}, its partitioning cost P_n^{cv}.

Note that the strategy "Clairvoyant" cannot be implemented algorithmically, since s and ℓ are not known until the classification is completed.

As shown in [AD15, Section 6], this strategy offers the smallest expected classification cost among all strategies that have oracle access to s and ℓ at the outset of a classification round. As such, its expected cost is a lower bound for the cost of all algorithmic classification procedures; hence we call it an *optimal strategy*.

The non-algorithmic strategy "Clairvoyant" can be turned into an algorithmic classification strategy, which is described next. It will turn out that its cost is only marginally larger than that of strategy "Clairvoyant".

Strategy "Count". *When classifying the ith element, for $1 \le i \le n-2$, proceed as follows: If $s_{i-1} \ge \ell_{i-1}$, compare with p first, otherwise compare with q first.*

The number of additional comparisons of this strategy is called A_n^{ct}, its cost P_n^{ct}.

No algorithmic solution for the classification problem can have cost smaller than "Clairvoyant". Strategy "Count" is algorithmic. Thus any cost-minimal algorithmic classification procedure has cost between $\mathbb{E}(P_n^{\mathrm{cv}})$ and $\mathbb{E}(P_n^{\mathrm{ct}})$, and a precise analysis of both will lead to good lower and upper bounds for the cost of such a procedure. It was shown in [AD15] that $\mathbb{E}(P_n^{\mathrm{ct}}) - \mathbb{E}(P_n^{\mathrm{cv}}) = O(\log n)$ and that, as a consequence, both strategies lead to dual-pivot quicksort algorithms that use $\frac{9}{5} n \log n + O(n)$ comparisons on average. In the following, we carry out a precise analysis of $\mathbb{E}(P_n^{\mathrm{cv}})$ and $\mathbb{E}(P_n^{\mathrm{ct}})$, which will make it possible to determine the expected comparison count of an optimal dual-pivot quicksort algorithm up to $0.28n$.

Lemma 9.1. *(a) The expected number of additional comparisons of strategy "Clairvoyant" is*

$$\mathbb{E}(A_n^{\mathrm{cv}}) = \frac{n}{6} - \frac{7}{12} + \frac{1}{4(n - [n \text{ even}])} - \mathbb{E}\big(X_n^{\searrow}\big)\,.$$

(b) The expected number of additional comparisons of strategy "Count" is

$$\mathbb{E}\big(A_n^{\mathrm{ct}}\big) = \frac{n}{6} - \frac{7}{12} + \frac{1}{4(n - [n \text{ even}])} + \mathbb{E}\big(X_n^{\nearrow}\big)\,.$$

Proof ideas: (The full proof can be found in Appendix J. A different proof of a related statement was given in [AD15].)

(a) Noticing that medium elements can be ignored, we consider a reduced input of size $n' = s + \ell$, consisting only of the s small and the ℓ large elements in the input. For $0 \le i \le n'$ let $s_i' = s - s_i$ and $\ell_i' = \ell - \ell_i$ denote the number of small respectively large elements left unclassified after step i. Then $\{(i, s_i' - \ell_i') \mid 0 \le i \le n'\}$ is a lattice path with distribution (including the distribution of n') exactly as in Section 7, so that the results on the expected number of zeros on such paths given there may be applied. We also note that the sign of $s_{i-1}' - \ell_{i-1}'$ decides whether the ith element to be classified is compared with p first or with q first, and that additional comparisons correspond to steps on the path that lead away from the horizontal axis, excepting down-from-zero steps (due to the asymmetry in treating the situation $s - s_i = \ell - \ell_i$ in strategy "Clairvoyant"). For the number of steps away from the horizontal axis one easily finds the expression $\min(s, \ell)$. Averaging over all choices for n' and the two pivots leads to the formula claimed in (a).

(b) Now assume strategy "Count" is applied to $n' = s + \ell$ elements. The set $\{(i, s_i - \ell_i) \mid 0 \le i \le n'\}$ forms a lattice path that starts at $(0, 0)$ and ends at $(n', s - \ell)$. It can be shown that reflection with respect to the vertical line at $n'/2$ maps these paths in a probability-preserving way to the paths from from (a) (and thus from our model), and it turns out that additional comparisons in this strategy correspond to steps away from the horizontal axis and up-to-zero steps. As in (a), averaging leads to the formula claimed in (b). □

Lemma 9.1 allows us to give an exact expression for the average number of comparisons of "Clairvoyant" and "Count" in a single partitioning step. The expressions for $\mathbb{E}(P_n^{\mathrm{cv}})$ and $\mathbb{E}(P_n^{\mathrm{cv}})$ are obtained by adding the expected number of necessary comparisons $\frac{4}{3}(n-2) + 1$ to the cost terms in Lemma 9.1 (see Appendix J).

10 Main Results and their Asymptotic Aspects

In this section we give precise formulations of our main results. We use the partitioning cost from the previous section to calculate the expected number of comparisons of the two dual-pivot quicksort variants obtained by using classification strategies "Clairvoyant" and "Count", respectively. We call these sorting algorithms "Clairvoyant" and "Count" again. Recall that "Clairvoyant" uses an oracle and is comparison-optimal, and that "Count" is its algorithmic version. We validated our main results in experiments which can be found in Appendix L. They show that the error term $O(n^{-4})$ is small already for real-life input sizes n, and that the linear term has a big influence even for larger n.

Theorem 10.1. *For $n \geq 4$, the average number of comparisons in the comparison-optimal dual-pivot quicksort algorithm "Clairvoyant" (with oracle) is*

$$\mathbb{E}(C_n^{\mathrm{cv}}) = \frac{9}{5}nH_n + \frac{1}{5}nH_n^{\mathrm{alt}} - \frac{89}{25}n + \frac{77}{40}H_n + \frac{3}{40}H_n^{\mathrm{alt}} + \frac{67}{800} - \frac{(-1)^n}{10} + r_n$$

where

$$r_n = \frac{[n \text{ even}]}{320}\Big(\frac{1}{n-3} + \frac{3}{n-1}\Big) - \frac{[n \text{ odd}]}{320}\Big(\frac{3}{n-2} + \frac{1}{n}\Big).$$

Corollary 10.2. *The average number of comparisons in the algorithm "Clairvoyant" is*

$$\mathbb{E}(C_n^{\mathrm{cv}}) = \frac{9}{5}n\log n + An + B\log n + C + \frac{D}{n} + \frac{E}{n^2} + \frac{F[n \text{ even}] + G}{n^3} + O\Big(\frac{1}{n^4}\Big)$$

with

$$A = \frac{9}{5}\gamma - \frac{1}{5}\log 2 - \frac{89}{25} = -2.6596412392892\ldots, \qquad B = \frac{77}{40} = 1.925,$$
$$C = \frac{77}{40}\gamma - \frac{3}{40}\log 2 + \frac{787}{800} = 2.042904116393455\ldots, \qquad D = \frac{13}{16} = 0.8125,$$
$$E = -\frac{77}{480} = -0.1604166\ldots, \qquad F = \frac{1}{8} = 0.125, \qquad G = -\frac{19}{400} = -0.0475,$$

asymptotically as n tends to infinity.

Before continuing with the second partitioning strategy, let us make a remark on the (non-)influence of the parity of n. It is noteworthy that in Corollary 10.2 no such influence is visible in the first six terms (down to $1/n^2$); only from $1/n^3$ on the parity of n appears. This is somewhat unexpected, since a term $(-1)^n$ appears in Theorem 10.1.

Theorem 10.3. *The average number of comparisons in the dual-pivot quicksort algorithm "Count" is*

$$\mathbb{E}\big(C_n^{\mathrm{ct}}\big) = \frac{9}{5}nH_n - \frac{1}{5}nH_n^{\mathrm{alt}} - \frac{89}{25}n + \frac{67}{40}H_n - \frac{3}{40}H_n^{\mathrm{alt}} - \frac{83}{800} + \frac{(-1)^n}{10} - r_n$$

where r_n is defined in Theorem 10.1.

Again, the asymptotic behavior follows from the exact result.

Corollary 10.4. *The average number of comparisons in the algorithm "Count" is*

$$\mathbb{E}\big(C_n^{\mathrm{ct}}\big) = \frac{9}{5} n \log n + An + B \log n + C + \frac{D}{n} + \frac{E}{n^2} + \frac{F[n \text{ even}] + G}{n^3} + O\Big(\frac{1}{n^4}\Big)$$

with

$$A = \frac{9}{5}\gamma + \frac{1}{5}\log 2 - \frac{89}{25} = -2.3823823670652\ldots, \qquad B = \frac{67}{40} = 1.675,$$
$$C = \frac{67}{40}\gamma + \frac{3}{40}\log 2 + \frac{637}{800} = 1.81507227725206\ldots, \qquad D = \frac{11}{16} = 0.6875,$$
$$E = -\frac{67}{480} = -0.1395833\ldots, \qquad F = -\frac{1}{8} = -0.125, \qquad G = \frac{31}{400} = 0.0775,$$

asymptotically as n tends to infinity.

The idea of the proofs of Theorems 10.1 and 10.3 is to translate the recurrence (8.1) into a second order differential equation for the generating function $C(z)$ of $\mathbb{E}(C_n)$ in terms of the generating function $P(z)$ of $\mathbb{E}(P_n)$. Integrating twice yields $C(z)$. This generating function then allows extraction of the exact expressions for $\mathbb{E}(C_n)$. The asymptotic results follow. See Appendix K for details.

Appendix

The appendices can be found at arXiv:1602.04031v1.

References

[AD15] Martin Aumüller and Martin Dietzfelbinger. Optimal partitioning for dual-pivot quicksort. *ACM Trans. Algorithms*, 12(2):18:1–18:36, November 2015.

[Dev16] The SageMath Developers. *SageMath Mathematics Software (Version 7.0)*, 2016. `http://www.sagemath.org`.

[GKP94] Ronald L. Graham, Donald E. Knuth, and Oren Patashnik. *Concrete mathematics. A foundation for computer science*. Addison-Wesley, second edition, 1994.

[Hen91] P. Hennequin. *Analyse en moyenne d'algorithmes : tri rapide et arbres de recherche*. PhD thesis, Ecole Politechnique, Palaiseau, 1991.

[HK15] Benjamin Hackl and Daniel Krenn. Asymptotic expansions in SageMath. `http://trac.sagemath.org/17601`, 2015. module in SageMath 6.10.beta2.

[Hoa62] Charles A. R. Hoare. Quicksort. *Comput. J.*, 5(1):10–15, 1962.

[Kra15] Christian Krattenthaler. Lattice path enumeration. In *Handbook of enumerative combinatorics*, Discrete Math. Appl. (Boca Raton), pages 589–678. CRC Press, Boca Raton, FL, 2015.

[Mah08] Hosam Mahmoud. *Polya Urn Models*. Chapman & Hall/CRC, 1 edition, 2008.

[Moh79] Sri Gopal Mohanty. *Lattice path counting and applications*. Academic Press [Harcourt Brace Jovanovich, Publishers], New York-London-Toronto, Ont., 1979. Probability and Mathematical Statistics.

[Sed75] Robert Sedgewick. *Quicksort*. PhD thesis, Standford University, 1975.

[WNN15] Sebastian Wild, Markus E. Nebel, and Ralph Neininger. Average case and distributional analysis of dual-pivot quicksort. *ACM Transactions on Algorithms*, 11(3):22, 2015.

[Yar09] Vladimir Yaroslavskiy. Replacement of quicksort in java.util.arrays with new dual-pivot quicksort. `http://permalink.gmane.org/gmane.comp.java.openjdk.core-libs.devel/2628`, 2009. Archived version of the discussion in the OpenJDK mailing list.